

A New Approach to Categorical Resampling

Daniel R. Steinwand¹
USGS EROS Data Center
SAIC
Sioux Falls, SD 57198
steinwand@usgs.gov

Abstract

A new method has been developed for resampling raster image data that contain class or categorical data. Categorical data are usually the result of an image classification or other statistical processes. During reprojection and resampling, the combination or interpolation of data with their neighboring pixels is not necessarily meaningful as it is with signal-based remote sensing data. The nearest neighbor resampling method is commonly used to resample this type of data. This method is chosen because the alternatives--cubic convolution, bilinear interpolation, etc. --are interpolating methods that do not preserve categorical values. In addition, the geometric distortions that are present in the projection change of data of global extent are far greater than distortions that occur in moderate- to high-resolution remote sensing data. Indeed, many of the software tools available today were designed for single-scene, signal-based remote sensing image data where the extent of the image is usually only a few hundred kilometers, rather than for datasets of global or continental coverage. The typical nearest neighbor resampling algorithm for categorical data takes into account only the center of the pixel and not the area covered by the pixel. In instances where the image (or a region of the image) is being undersampled, nearest neighbor resampling can result in imagery that is not representative of the original image. The new resampling method treats pixels not as points, but as areas. This approach maps the corner coordinates of the output image pixel back into the input image and statistically determines the pixel value on the basis of input image pixels that lie within the output pixel's geometric extent.

¹ USGS EROS Data Center, SAIC, Sioux Falls, SD 57198-0001. Work performed under U.S. Geological Survey contract 03CRCN0001.

Introduction

In recent years, advances in remote sensing technologies and computer storage and processing power have enabled the construction of many datasets containing raster image data of global or continental extent. These datasets are typically combinations of several data acquisitions and are often processed beyond the raw image sensor data. A common data representation is the composite pixel (Eidenshink and Faundeen, 1994) that is a representative value of a particular ground location over a period of time. Also present are classification images or images produced with other statistical processes. The pixels in these types of data are no longer signal-based data but are statistically binned into classes or categories. These data are often referred to as “class” or “categorical” data; each pixel is assigned a value to represent a given class for the portion of the Earth that it covers. A common example of this is land cover classification data, as contained in the Global Land Cover Characterization Database developed by the U.S. Geological Survey's Earth Resources Observation System Data Center and others (Loveland, et al., 1999).

The construction of these categorical global datasets is often a very large task. As computational power and computer storage technologies advance, so do the requirements of the scientific community; the size of these datasets and the amount of processing incurred to create them are enormous (Eidenshink et al, 1994, Steinwand, 1994). In addition, because these datasets are often combinations of many different acquisition dates, the choice of a single projection is necessary for production purposes and is often chosen on the basis of processing and storage considerations rather than the eventual use of the data. As a result, data are often offered in one projection, and a single projection rarely fills the needs of all end users.

End users of these data often find it necessary to reproject the data over a given study area so that they can be combined with data from other sensors and studies. Very often, those studying the Earth at a global scale will require a different projection when studying a single continent. Therefore reprojection, and thus, resampling, are necessary.

The Problem With Reprojection and Resampling

Reprojection and resampling categorical raster image data of global or continental extent should be approached with caution. Users of these data are often presented with the following situation: They need to present or perform their study in a given projection over a given study area. Often the data needed to perform this study exist in a different projection and pixel size. Most modern image processing or geographic information system software will perform these transformations, but the issues of geometric distortion and the errors due to resampling need to be considered carefully.

The errors due to geometric distortion that occur during a projection change of raster image data have recently been documented (Steinwand, et al., 1995, Mulcahy, 2000,

Mulcahy and Clarke, 2001, Usery and Seong, 2001). The errors due to resampling in areas of large geometric distortion or scale change caused by a change in projection have not yet been adequately addressed.

The nearest neighbor resampling algorithm computed by means of the inverse mapping process takes one point in the output space pixel (usually the pixel center as presented here, but some systems standardize on one of the pixel corners) and maps that point into the input image space (Figure 1). This transformation usually consists of a linear transformation from the output space line and sample coordinate to the rectangular projection coordinate of the output space projection. This coordinate is then reprojected into the projection system of the input image space. Another linear transformation is then applied to get to the input space line and sample coordinate. This resulting input space image coordinate is usually not at an exact pixel location, so it is rounded to the nearest pixel position, and that pixel's value is assigned to the output image at the output image coordinate under study. Although this method is computationally efficient, it can result in imagery that is not representative of the original data.

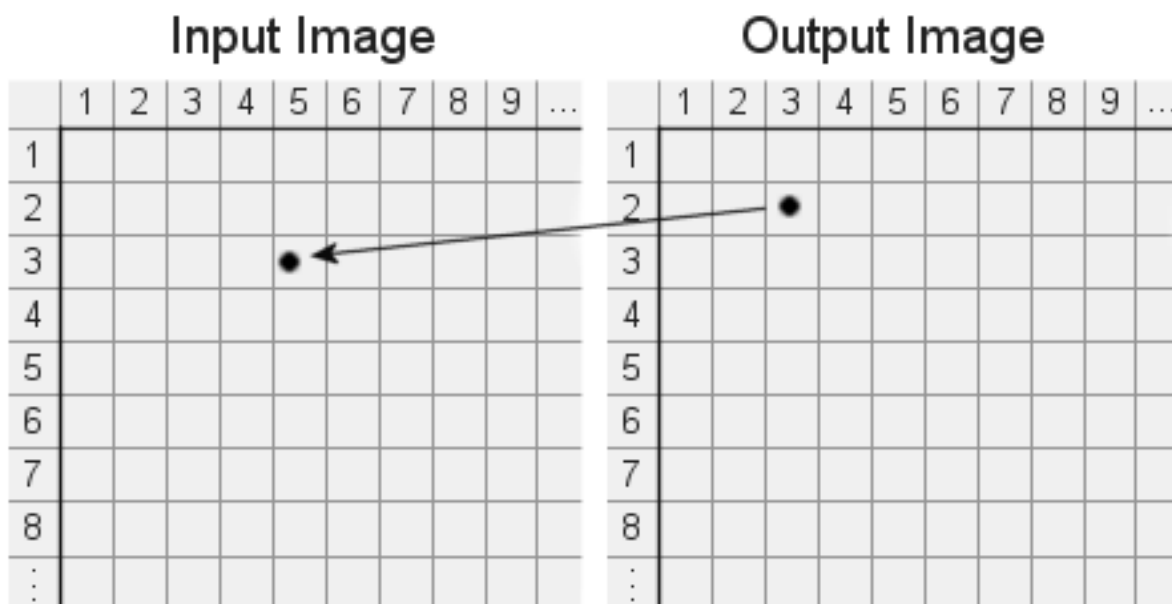


Figure 1. Nearest neighbor mapping of 1 pixel.

When the output imagery is constructed so that the resolution of the image is reduced, or when the projection change transformation causes this situation to occur, the next pixel in the output image, again mapped with the same algorithm, falls more than 1 pixel away in the input image (Figure 2). From a signal theory point of view, the input imagery in this case is undersampled.

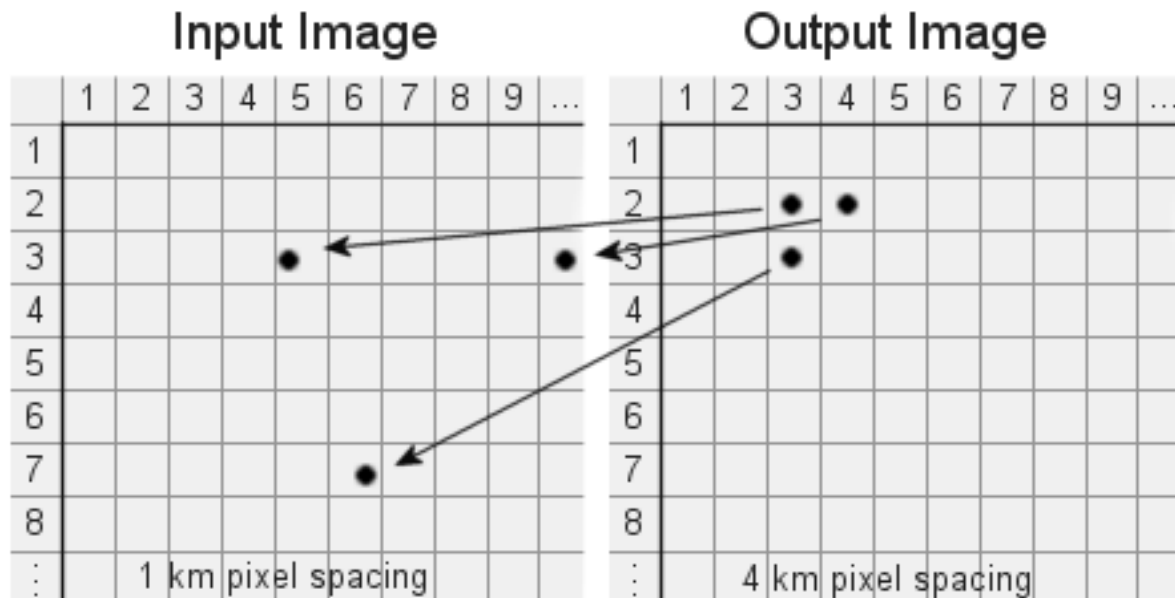


Figure 2. Nearest neighbor mapping of adjacent pixels.

Under these conditions, not all data in the input space image are used in the output image. More importantly, the nearest neighbor resampling algorithm does not necessarily choose a pixel that is representative of the area being resampled; it just takes the nearest one. This can result in parts of an image dataset not being representative of the input image area if, for example, a minority class happened to be the nearest pixel. Another way to think of this (in the scale of Figure 2) is to state that we have 1-km pixels spaced 4-km apart in the output image instead of a pixel that truly represents the area of the 4-km pixel.

The New Algorithm

The new resampling algorithm treats pixels not as points, but as areas. The four corners of each pixel (instead of just the center) are mapped into the input space. As shown in Figure 3, the corners (*a*, *b*, *c*, and *d*) of the output pixel at sample 3, line 2, map to input locations A, B, C, and D. From this, we see that 23 pixels (some of them partially—see below) make up the data that could be considered for the output pixel. The new algorithm uses the coordinates at points A, B, C, and D and creates a polygon. Any pixel that has its center point within the polygon is considered to be inside the area and is included in the statistical resampling.

Once the input image pixels that fall under the output image pixel footprint are determined, simple statistical methods, such as the maximum occurring pixel, the minimum, the average, etc., can be applied to determine the output image pixel value that is to be assigned to the output image. More complex methods could also be used (at the cost of

runtime performance) that favor certain classes or that combine classes into composite classes.

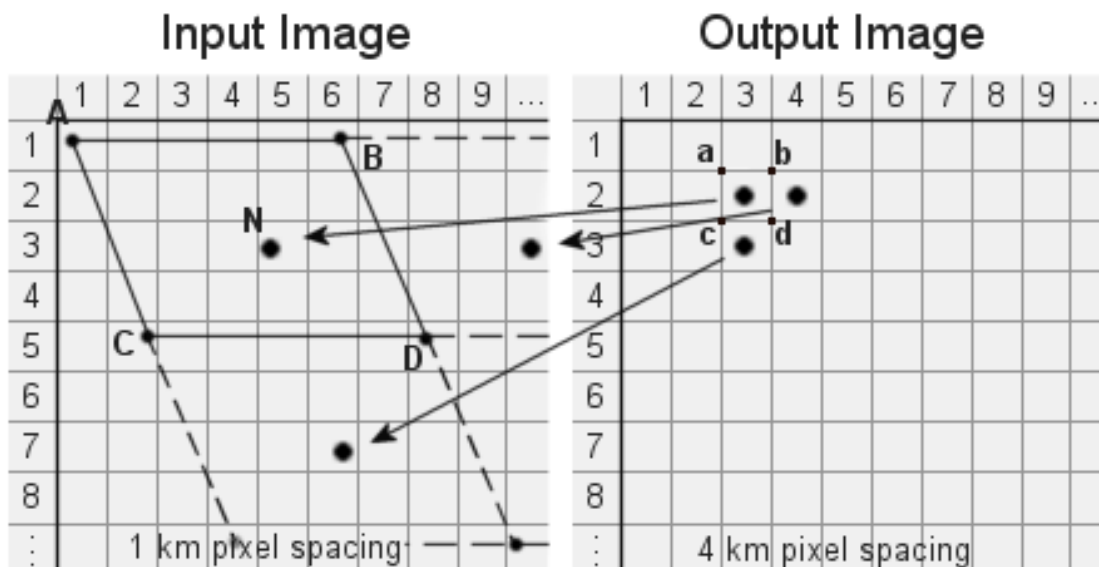


Figure 3. Mapping the input pixel's footprint—the concept behind the new algorithm.

An Example

The following example illustrates how the algorithm works for a 3-pixel by 2-line output area for the case where the output imagery has been projected to the Mollweide projection with a 4-km pixel size. The original imagery is in the Plate Carrée projection (geographic coordinates) with an approximate pixel size of 1-km. Because the original imagery is nearly 1 Gb, this type of subsampling (four-to-one) is not uncommon and illustrates the need for the new resampling method.

The algorithm follows the method of inverse mapping defined above. For each pixel, five coordinates are mapped: The pixel's center is mapped to determine the nearest neighbor mapping coordinate (this can be omitted in a production algorithm), and each of the output pixel's four corners are also mapped to the input image space. Note that these resulting input image coordinates are often not integer locations, and as such they do not map to exact pixel locations; they fall between pixels. Figure 4 shows a part of the Global Landcover Characteristics Image over south Florida in Plate Carrée, and Figure 5 shows the results of the resampling and reprojection process (5a with nearest neighbor and 5b with the new method). Figure 6 shows a magnification of a small area near the center of Figure 5, an area just to the east of Tampa, Fla. Figure 6a shows the results of nearest neighbor resampling, and Figure 6b shows the results of the new algorithm. The remainder of this example focuses on the 3- by 2-pixel region (circled) and the differences in results between the two methods.

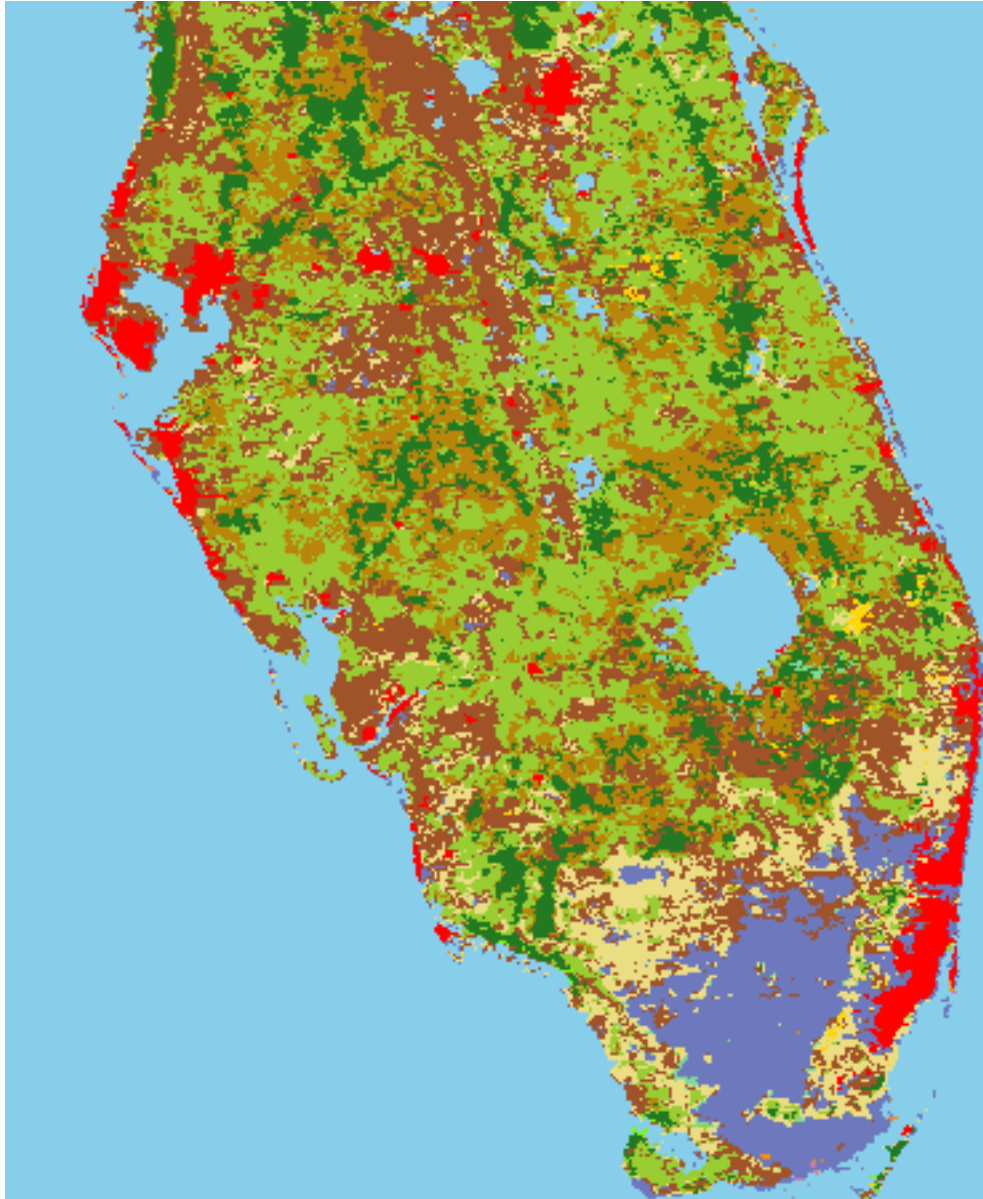


Figure 4. A part of the input image.

Figure 7 shows the part of the input image that corresponds to the circled area in Figure 6; it has been enlarged to show the footprint of the individual output space pixels that go into making the output image. (Note that the area labeled *a* has the same geometry as used in Figures 1 - 3.) The nearest neighbor (center pixel of each area) is marked with a solid black dot. This is the pixel that is assigned to the output image when the nearest neighbor resampling method is used. The four corners of the output image pixel are also drawn, with lines connecting them to show the footprint of the output pixel in the input image space.

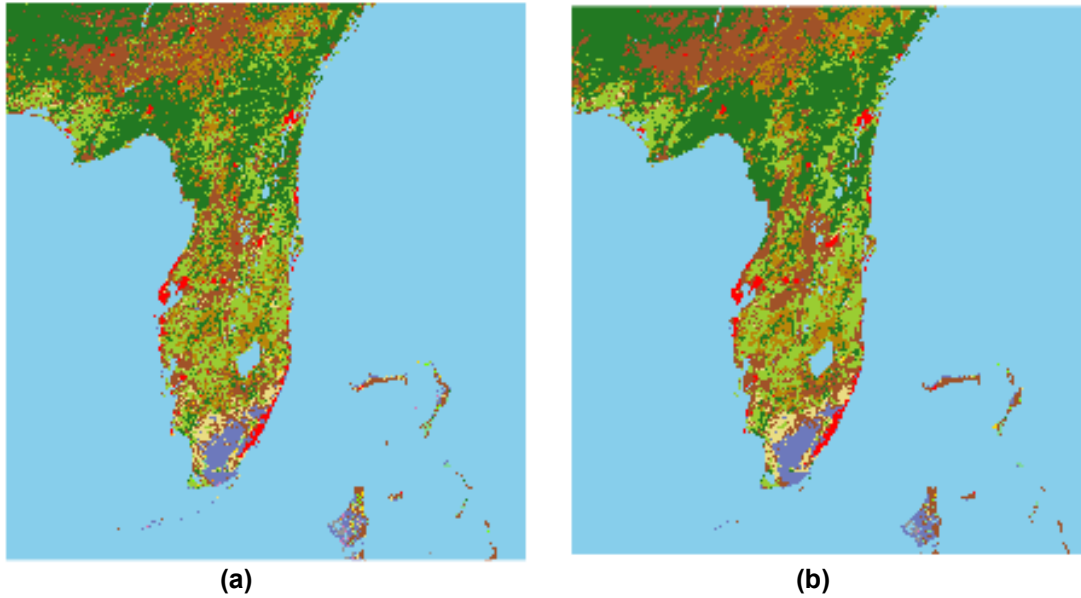


Figure 5. Results of reprojection to Mollweide at 4-km pixels. Nearest neighbor results are on the left (a); and results of the new algorithm are on the right (b).

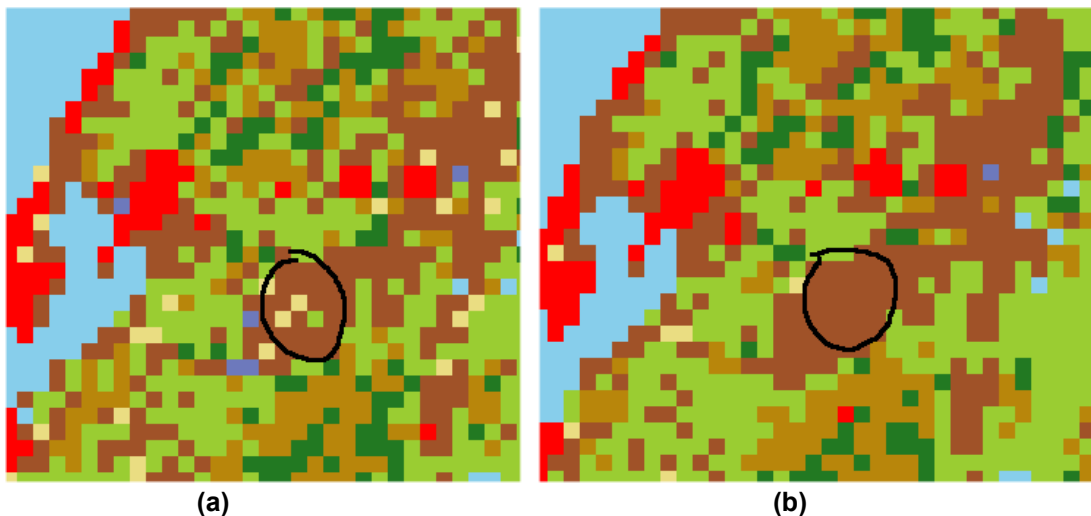


Figure 6. Enlargement of the center part of Figure 5, with the study area circled.

When nearest neighbor resampling is applied, pixels *a*, *c*, and *e* are assigned the *Dryland Cropland & Pasture* class, pixels *b* and *d* are assigned the *Grassland* class, and pixel *f* is assigned the *Cropland/Woodland Mosaic* class. When the new algorithm is applied using the maximum occurring pixel method, pixels *a*, *b*, *c*, *d*, *e*, and *f* are all assigned the *Dryland Cropland & Pasture* class--the class that occurs the most often in each of the pixel areas shown. A closer look at the pixels with results that differ between the two methods, pixels *b*, *d*, and *f*, follows.

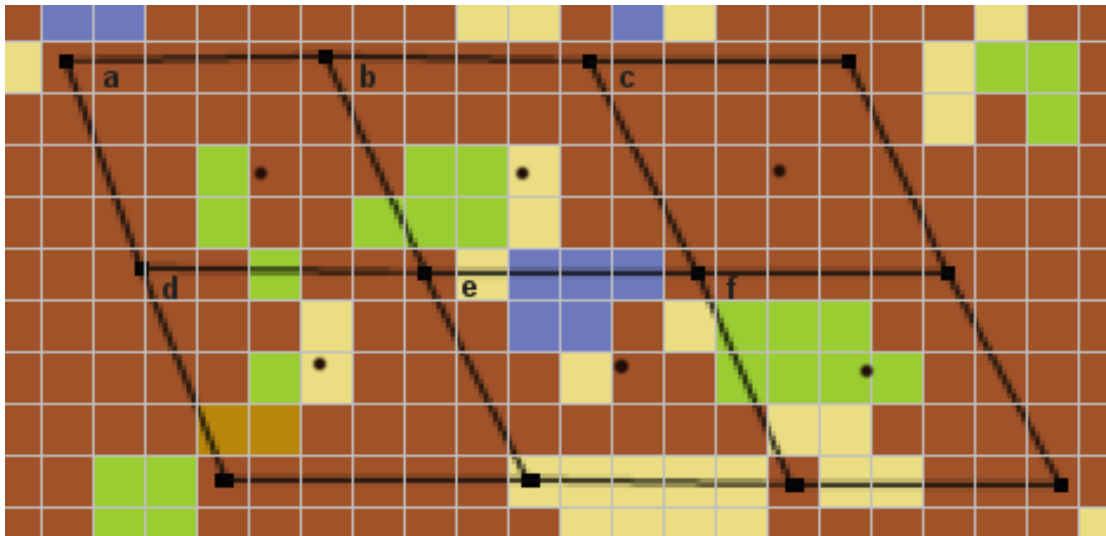


Figure 7. Enlargement of input area showing individual output pixel footprints. The brown pixels represent the *Dryland Cropland & Pasture* class, green the *Cropland/Woodland Mosaic* class, yellow-tan the *Grassland* class, blue the *Wooded Wetland* class, and dark-tan the *Cropland/Grassland Mosaic* class.

In pixel *b*, the nearest neighbor pixel is the *Grassland* class--represented by only two pixels of the 20 (10%) that go into the area that makes up output pixel *b*. Four pixels of the 20 (20%) represent the *Cropland/Woodland Mosaic* class, and the remaining 14 pixels (70%) represent the *Dryland Cropland & Pasture* class, which was chosen by the new algorithm. A similar situation occurs with pixel *d*.

In pixel *f*, the maximum occurring class is the *Dryland Cropland & Pasture* class, with 11 of 20 pixels (55%) belonging to that class. The *Cropland/Woodland Mosaic* class, which is also the class chosen by the nearest neighbor methods, has 7 of the 20 (35%), and the *Grassland* class has 2 of 20 (10%). In this situation, it may be more correct to produce some sort of combination of the two dominant classes, but the added complexity could increase runtime and make the interpretation of results more complex.

When the geometric characteristics of the projection change map to a single pixel, or, in the case of an enlargement, to a fraction of a pixel, the algorithm reverts back to the nearest neighbor method of resampling.

Visual Results

Visually, the results of processing with the new algorithm are most apparent during extreme downsampling and reprojection conditions. This is common for publishing pictures in reports and on the Web. The two images below illustrate the output of the algorithm. Figure 8 was processed using the nearest neighbor resampling method and appears noisy. Figure 9 was processed with the maximum occurring pixel method and

appears smoother. Both figures 8 and 9 incorporate a reprojection to Mollweide and a downsampling to 50-km pixels from the original 1-km data.

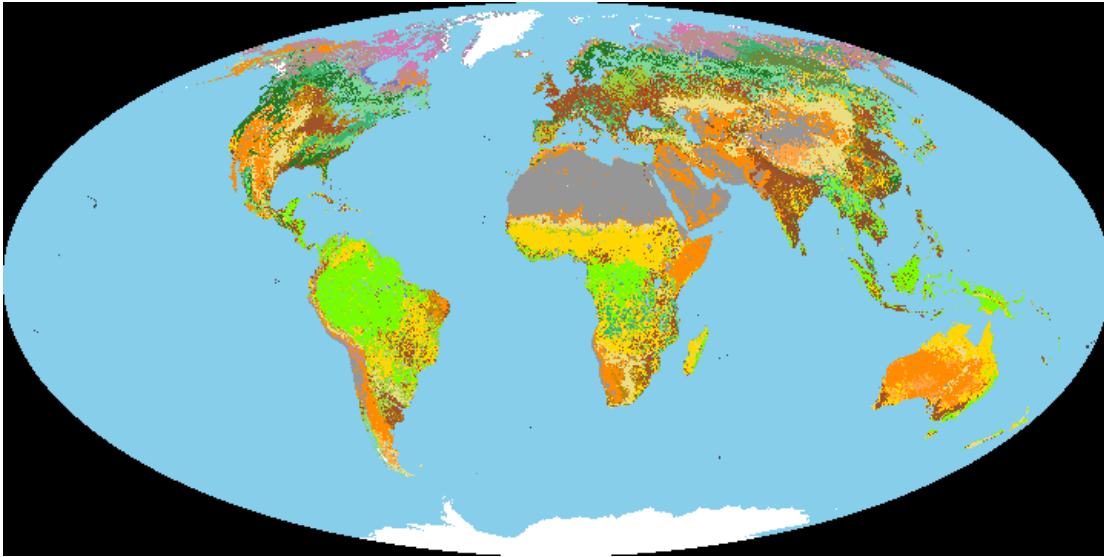


Figure 8. Extreme downsampling and reprojection with nearest neighbor.

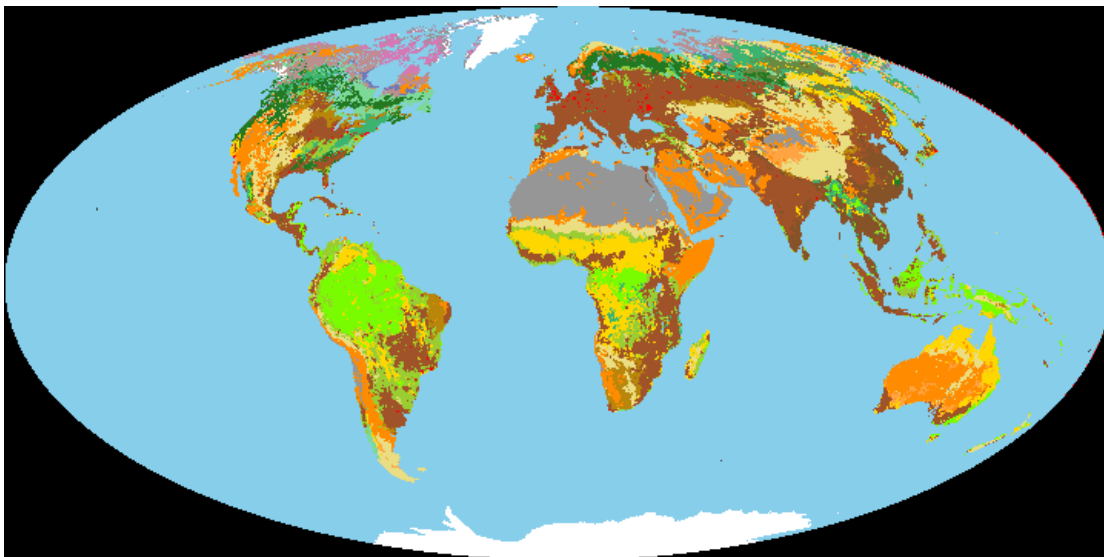


Figure 9. Extreme downsampling and reprojection with the new algorithm.

Conclusion

A new resampling method has been presented that takes into account pixel areas instead of merely a point within a given pixel. The algorithm maps the footprint of the output pixel back into the input image and takes into consideration all pixels under that footprint, rather

than just the nearest neighbor. The current algorithm includes the following statistical resampling methods: The maximum occurring pixel, the minimum occurring pixel, and the nearest neighbor. There is also an option to include a most preferred class and to exclude a least preferred class with both the minimum and maximum methods. Other statistical methods, perhaps those that take into account spatial patterns, combination classes, or class distribution, could be implemented within the existing algorithm framework.

References

Eidenshink, J.C., and Faundeen, J.L., 1994, The 1-km AVHRR global land data set: First stages in implementation: *International Journal of Remote Sensing*, Taylor and Francis, Ltd., v. 15, no. 17, p. 3443-3462.

Eidenshink, J.C., Steinwand, D.R., Wivell, C.E., Hollaren, D.M., and Meyer, D.J., 1994, Processing techniques for global land 1-km AVHRR data, in Pecora 12 Symposium, Land Information from Space Based Systems, Sioux Falls, South Dakota, August 1993, Proceedings: Falls Church, Virginia, American Society for Photogrammetry and Remote Sensing, p. 214-222.

Loveland, T.R., Zhu, Zhiliang, Ohlen, D.O., Brown, J.F., Reed, B.C., and Yang, Limin, 1999, An analysis of the IGBP global land cover characterization process: *Photogrammetric Engineering and Remote Sensing*, v. 65, no. 9, p. 1021-1032.

Mulcahy, K.A., 2000, Two new metrics for evaluating pixel-based change in data sets of global extent due to projection transformation: *Cartographica*, v. 37, no. 2, summer 2000, p. 1-11.

Mulcahy, K.A. and Clarke, K.C., 2001, Symbolization of map projection distortion: A review: *Cartography and Geographic Information Science*, v. 28, no. 3, p. 167-181.

Steinwand, D.R., 1994, Mapping raster imagery to the Interrupted Goode Homolosine projection: *International Journal of Remote Sensing*, Taylor and Francis, Ltd., v. 15, no. 17, p. 3463-3471.

Steinwand, D.R., Hutchinson, J.A., and Snyder, J.P., 1995, Map projections for global and continental data sets and an analysis of pixel distortion caused by reprojection: *Photogrammetric Engineering and Remote Sensing*, Falls Church, Virginia, American Society for Photogrammetry and Remote Sensing, v. 61, no. 12, December 1995, p. 1487-1497.

Usery, E.L. and Seong, J.C., 2001, All equal-area map projections are created equal, but some are more equal than others: *Cartography and Geographic Information Science*, v. 28, no. 3, p. 183-193.